

2007

The IRS Research Bulletin

Publication 1500

The IRS Research Bulletin

Proceedings of the 2007 IRS Research Conference

Research
Analysis
Statistics

A Framework for Optimal Tax Administration

Alan H. Plumley, Internal Revenue Service

Few people concern themselves with how well the tax system is administered—until they have to deal with the tax administrator. Worse yet, there is very little consensus on how the tax system *should* be administered—particularly what its ultimate objective should be. Fortunately, more and more people are asking the right kinds of questions. For example, in connection with recent hearings on the tax gap, members of Congress asked the Internal Revenue Service (IRS) several important questions:

- What is the right-sized IRS budget to increase voluntary compliance by 1 percent? By 5 percent?
- How would the rate of voluntary tax compliance be affected by sustained growth in the IRS budget of \$500 million over each of 5 years?
- What would be the optimal use of these resources in providing improved services, stronger enforcement, and enhanced information technology?

Unfortunately, the answer to each of these questions is that we do not currently know. Therefore, this paper attempts to provide a framework for how a tax administration agency might address these, and similar, questions. I need to emphasize, however, that our understanding of what constitutes optimal tax administration is not very advanced. What follows, then, is not an official IRS position, but rather one observer's attempt to move the debate along by framing important questions and by suggesting various ways of addressing them. This approach uses a benefit-cost framework that is very familiar to economists and many policymakers. In the context of this framework, I lay out a series of principles and steps that the IRS could decide to take to improve tax administration. These steps may not all be feasible, but the issues they are aimed at will need to be addressed if the IRS is to make significant strides toward improved tax administration. In order to answer such crucial questions, the IRS and its stakeholders need to answer three more fundamental questions:

1. **Theory:** How can we identify the best use of resources to achieve our objective, given our constraints?

2. **Objective:** What, specifically, should the IRS be trying to achieve?
3. **Plan:** What should we do in the short and long term to make progress toward optimal use of IRS resources?

Although there is much room for debate on these questions,¹ there are 15 key principles that I believe should guide the quest for optimal tax administration. I list and explain those principles below, organized (except for the first, introductory, one) under those three fundamental questions. My focus is on Federal tax administration in the U.S., but most of the principles probably apply equally well to other tax administrations as well. I again emphasize that these are primarily my own views, rather than those of the Office of Research more generally. This paper is meant to stimulate thinking and discussion about optimal tax administration.

Principle #1. The stakes are huge. Suboptimal tax administration results in too little voluntary compliance (e.g., the gross tax gap is estimated to be \$345 billion for Tax Year 2001), a large number of inadvertent errors by taxpayers, unnecessary compliance burden placed on taxpayers, and a disproportionate share of the overall tax burden borne by compliant taxpayers. Furthermore, not having a framework for optimal tax administration makes it impossible to know how to manage the IRS budget as effectively as possible. This gets right to the heart of what the IRS should be doing—affecting not just budget deficits, but also touching the lives of most Americans very directly.

Theory: How Can We Identify the Best Use of Our Resources?

Before we can achieve optimal tax administration, we need to define it. Fortunately, there is a well-established body of theoretical thought that can be applied to tax administration. This gives rise to the next five principles.

Principle #2. Optimal tax administration can be modeled as a constrained optimization problem, with allowance for off-model judgments when necessary. The administration of tax laws can be guided by a single, explicit objective, subject to explicit constraints. The objective function specifies what is to be optimized. In the case of tax administration (as in many other contexts), the objective ought to be to maximize the net benefits (i.e., benefits minus costs) of all our activities (i.e., our uses of resources).

¹ See Plumley and Steuerle (2004).

The constraints account for budgetary and other practical limitations, which may or may not vary over time. Formally stated, the ultimate objective of the IRS is to:

Maximize: Benefits – Costs

Subject to: Budget and other constraints

Defining the benefits, costs, and other constraints, of course, is critical and the subject of the next section (“Ultimate Objective”). However, before considering those details, four additional principles should be emphasized.

Principle #3. Model features that cannot (yet) be quantified empirically or theoretically can in some cases be quantified (or at least plausibly bounded) by consensus assumptions. Otherwise, such factors should be accounted for subjectively to modify the model’s output. The benefit of deriving explicit assumptions whenever possible by developing consensus among the key decisionmakers is that the ones who ultimately make the decisions can thereby establish rules for themselves from a holistic (rather than parochial) perspective in an internally consistent, objective (rather than crisis) atmosphere. Not (yet) being able to quantify certain components of the model empirically must not cause us to abandon the quest to do so, nor prevent us from becoming progressively more rigorous over time.

Principle #4. The optimal solution to the constrained optimization problem is the allocation of resources that equalizes the marginal benefit/cost (“bang for the buck”) across all opportunities to use those resources, as long as the benefits and costs include all of the elements of the ultimate objective. If the marginal benefit/cost ratios (i.e., the benefit produced by the last dollar or the next dollar spent) are not equalized, then more net benefit can be obtained by shifting resources from activities with low marginal benefit/cost ratios to those with higher benefit/cost at the margin.

Principle #5. Whether allocating the budget across organizations and programs, selecting discretionary workload within programs, or selecting discretionary issues to pursue on a given case, the same ultimate objective should be applied. However, not all components of the ultimate objective will necessarily apply in each of these contexts. For example, the impact of an audit on the voluntary compliance of the general population may not be affected by how many lines on the tax return are examined. The principle that there should be just one ultimate objective means, for example, not only that, when resources are applied optimally, the benefit arising from the last dollar spent on audits of large corporations should equal the benefit of the last dollar spent on the document matching program, the benefit of

the last dollar spent on taxpayer assistance, and the benefit of the last dollar spent on criminal investigations, but also that the same agencywide marginal benefit/cost ratio governs both which enforcement cases to pursue and which issues to pursue on those cases. That is, cases and issues with an expected marginal benefit/cost ratio less than the agencywide ratio should not be pursued. It also means that, if the budget were expanded, the agencywide marginal benefit/cost ratio would be decreased, enabling more cases to be worked (at the extensive margin) and more discretionary issues to be pursued (at the intensive margin) on all cases worked—even on those cases that would have been worked with a smaller budget.

Principle #6. Other (more common) approaches to tax administration almost certainly do not lead to optimal decisions, will often lead to internally inconsistent decisions, and should be avoided. These include allocating resources and selecting workload based on:

- **The distribution of the tax gap:** Although it may seem natural to focus efforts on the areas where noncompliance is worst, that is generally not optimal for operational programs. That is because the largest components of the tax gap are often also the components that are the least cost-effective to combat (at least with current methods). That is often why noncompliance is worse among those components. Ultimately, expanding current activities in those components would typically result in less of the tax gap being reduced than if the same resources were applied to more cost-effective activities.
- **The noncompliance rate in each activity:** This is really the same as the previous point, since the noncompliance rate is merely the tax gap divided by the amount of true tax liability.
- **Average benefit/cost in each activity:** This presumes that all additional work in a given activity will be, on average, just as cost-effective as work in the “base.” That would be the best assumption if we selected workload completely randomly, but for most types of work, we attempt (and generally manage) to select the most cost-effective workload within each activity, resulting in declining marginal cost-effectiveness as the level of effort increases.
- **Direct yield in each activity:** This is quite common both for workload selection and in allocating resources across various activities—partly because it is easy. Unfortunately, it typically

uses a measure of yield other than the tax that is actually paid, and it focuses on direct enforcement results only. Moreover, the cost necessary to generate the yield is often not taken into account in a rigorous way, if at all. As a result, the overall net benefit generated is bound to be significantly less than could have been achieved with a more optimal allocation of the same resources.

Even if all we were interested in were maximizing direct enforcement dollars, as illustrated in the example below, we must be guided by cost-effectiveness, not the average. In this example, cases 2 and 3 each (and in combination) have a lower average yield than case 1, but, together, they generate more yield than case 1 for the same cost and should, therefore, be pursued before case 1.

Table 1. Hypothetical Comparison of Yield/Cost and Average Yield

Case	Expected yield (A)	Expected cost (B)	Yield/cost (A/B)		Average yield
1	200	100	2.0		200
2	150	50	3.0	2.75	137.5
3	125	50	2.5		

- No-change rates:** If a taxpayer is audited, for example, and the audit determines that the taxpayer reported the correct amount of tax, that audit is generally thought of as a “waste” of both the taxpayer’s and the IRS’s time and efforts. Therefore, IRS has historically sought to minimize the number of such cases as a percentage of all cases worked. There is a problem with using that as the sole objective, however: it is quite possible that a category of work with a relatively high no-change rate nonetheless has a very high overall marginal benefit/cost ratio—even after taking into account the fact that the cost incurred on the no-change cases yields no direct benefit. That is because the direct benefit of the change cases and the indirect effects of both the change and the no-change cases can compensate for the lack of direct benefit from the no-change cases. Without quantifying all of those, it is not clear that we should minimize no-change rates. In fact, once we quantify those component effects, we would be able to allocate resources according to marginal benefit/cost, so that the no-change rate would be unnecessary anyway.

- **Rules of thumb:** Ratios of amounts reported on a return, average amounts reported, and comparisons with prior years or with other taxpayers, for example, can be very misleading. Although such rules of thumb can be helpful in some cases, they invariably do not take everything into account, and they are not a good substitute for rigorous estimates of benefit/cost that take into account all the relevant factors. Unless a rule of thumb is itself a rough indicator of benefit-to-cost, it will not likely be a helpful basis for allocating resources.
- **Vague perceptions of noncompliance:** Lacking specific empirical findings, it is tempting to make decisions on the basis of subjective perceptions. Yet, unless these subjective judgments attempt to quantify marginal benefit/cost, and are all that is available in the short term, they could easily do more harm than good.

Ultimate Objective: What Should We Be Trying to Achieve?

Principle #7. The current IRS mission statement is not specific enough to identify, by itself, the optimal way to administer the tax laws. It outlines general purposes and values only.

The current (1998) IRS Mission Statement is:

Provide America's taxpayers top quality service by helping them understand and meet their tax responsibilities and by applying the tax law with integrity and fairness to all.

In contrast, the previous (1984) IRS Mission Statement was:

The purpose of the Internal Revenue Service is to collect the proper amount of tax revenue at the least cost to the public, and in a manner that warrants the highest degree of public confidence in our integrity, efficiency, and fairness.

This statement was closer to a statement of ultimate objective and can serve as a guide for resource allocation. For example, it identifies “collect[ing] the proper amount of tax” as the primary benefit to be sought. This includes both voluntary and enforced payments. “Proper” acknowledges that taxpayers should be refunded any overpayments just as much as they should be required to pay any underpayments. “Collect the proper amount of tax” also implies two more things: what really matters is what is

finally paid (not just what is reported on returns or recommended or assessed through enforcement), and it is tax (not interest and penalties) that is the ultimate benefit. The statement also makes it clear that these benefits should be derived as cost-effectively and fairly as possible. Put another way, the costs borne by the public (both IRS costs and taxpayer burden), and the steps that the IRS takes to foster “the highest degree of public confidence in our integrity, efficiency, and fairness,” are all means to the end of collecting the proper amount of tax; they are not ends in and of themselves. Although not everything that the IRS must do has a tax revenue impact (either directly or indirectly), the 1984 Mission Statement captures the bulk of IRS’s role in administering the tax laws. A more explicit statement of this objective follows in Principle #8.

Principle #8. The ultimate objective of the IRS is to maximize the weighted net benefits (weighted benefits minus weighted costs), subject to the relevant constraints. Specifically, the benefits, costs, weights, and constraints should include the following:

Benefits: There are six major benefits that IRS should seek to achieve: tax revenue paid timely, enforcement revenue paid sooner rather than later, other late payments, revenue protected, overpayments refunded through efficient returns processing and enforcement, and tax payments more closely aligned with true liabilities. These are explained below.

Tax revenue paid timely: This is the ultimate intent of the Tax Code and the ultimate benefit (both to the IRS and to taxpayers). In practice, it encompasses several things:

- Tax paid in regular deposits throughout the year, both directly by taxpayers and indirectly through withholding agents;
- Tax remitted with filed returns;
- The indirect effect (i.e., improved voluntary compliance in the general population) of taxpayer service activities; and
- The indirect effect of enforcement activities. This is generally considered the deterrent effect of IRS enforcement. It includes both the improvement in subsequent compliance behavior among those who are contacted through enforcement (the “subsequent-year effect”) and the improvement in the voluntary compliance of those who are *not* contacted (the “ripple effect” of enforcement). In reality, it is quite likely that the ripple effect arises from both a deterrent mechanism

and an assurance mechanism. The latter is likely the explanation for an improvement in compliance among those who do not consider themselves probable targets of a particular type of enforcement, such as criminal investigations; they probably improve their compliance behavior (or do not erode it) when they conclude that scofflaws are being found and punished fairly and effectively, giving them greater assurance that the tax laws are being administered fairly and taking away potential reasons to participate in petty noncompliance in protest themselves.

Enforcement revenue paid sooner rather than later: It would certainly be better (e.g., less costly and quicker) if all tax obligations were paid voluntarily and timely, but, if payments have to be enforced, those tax payments are nonetheless benefits that should be taken into account. Specifically:

- This is the direct effect of enforcement activities (additional revenue paid late because of enforcement contacts).
- The actual benefit is dollars of tax collected (not just recommended or assessed).
- Penalties and interest should probably be weighted less than tax (possibly not treated as benefits at all for resource allocation purposes). Even if the inclusion of interest fully accounts for the time value of the money, there are added benefits to the money being paid on time. That is, we should not prefer an outcome that imposes penalties on taxpayers (and therefore increases revenues) over one in which they pay in full and on time.
- Late payments of tax should similarly be weighted less than timely payments. This could be handled by applying an agreed-upon discount rate to late payments.

Other late payments: These are amounts of tax that are paid late, but they are paid without IRS intervention. Examples include:

- Amounts paid late because the taxpayer forgot to enclose a check for the balance due with the return. As long as the payment is made before the IRS sends the taxpayer a notice, it is not considered enforcement revenue.

- Amounts paid before a filing extension but after the payment due date (which is generally the original due date of the return).
- Amounts paid with amended returns that are filed voluntarily after the original due date.
- Amounts paid with returns that are filed late without any IRS prompting.
- Amounts paid in response to IRS soft notices (e.g., attempts to clarify information or suggest future behavior, without actually making an assessment).

As with enforced payments, these other late payments should be weighted less than timely payments in the ultimate objective function.

Revenue protected: Money that would have been paid out to taxpayers (but should not have been), were it not for IRS intervention, should be counted as a benefit of IRS activities. However, whenever we attempt to estimate total benefits, to avoid double-counting, we should include only amounts that are not already included with timely paid tax, such as disallowed refundable credits and claims for refunds on fraudulent returns.

Overpayments refunded through efficient returns processing and enforcement: We should not include overpayments as benefits; they are paid unnecessarily or by mistake. Whenever we attempt to estimate total benefits, they should be netted out. However, the amount of each valid refund should be considered a benefit for resource allocation and workload selection purposes. This should include:

- Valid refunds claimed on timely or amended returns;
- Overpayments detected by math error and returns processing; and
- Overpayments detected by enforcement.

It is important to note, however, that while amounts that taxpayers overpaid and were refunded should be considered as positive benefits (rather than subtracted from total benefits), overpayments that neither the taxpayers nor the IRS detect should not be consid-

ered a benefit. Rather, they are a cost—most likely attributable to complexity, ignorance, confusion, etc. Therefore, undetected overpayments should be considered a type of compliance cost (or burden). See “Taxpayer and Other Private Costs” below.

Having tax payments more closely aligned with true liabilities:

Just as compliant taxpayers benefit from the fairness achieved when their overpayments are corrected, they also benefit when noncompliant taxpayers are required to pay their tax in full. This sometimes has monetary benefits that are not already taken into account in the refunds and enforced payments described above. For example, tax enforcement generally fosters fair business competition, which allows compliant taxpayers to stay in business and earn higher profits. Moreover, compliant taxpayers generally benefit from a smaller tax gap to the extent that it is associated with some combination of lower tax burdens, lower budget deficits, or more public goods and services.

Costs: There are two main categories of cost. These are generally the means to achieving the benefits. Marginal benefit/cost estimates that are used for resource allocation must reflect the full cost of the activities in question—including all overhead, indirect, and follow-on costs.

IRS Costs: These are the amounts included in the IRS budget appropriations for the year in question. For any current fiscal year, these amounts are fixed. IRS can perhaps influence the budget somewhat for future years but not in the current year. Resources are therefore limited, forcing a prioritization of work—identifying the best combination of tax administration activities and the best workload to maximize net benefits. Having separate budget appropriations for different IRS activities imposes some constraints on this optimal use of the overall level of resources, however, and these constraints may or may not significantly limit the overall net benefits that can be achieved.

Taxpayer and Other Private Costs: Taxpayers incur various costs in order to comply with their tax obligations. These take the form of out-of-pocket expenses, time, and frustration. Taxpayers need to keep records, obtain forms and publications, complete and submit forms, make payments, etc. Many taxpayers choose to pay for goods and services that will help them with these tasks. In addition, third parties incur costs to provide information reporting and withholding services for taxpayers and the IRS. All of these

costs should be taken into account when IRS allocates its resources, though IRS costs and private costs may be weighted (valued) differently. (See Principle #11 for more details.)

Weights: When components of Benefits and Costs are combined, they will likely need to be weighted differently to reflect differences in inherent value. For example:

- Late payments should be weighted less than timely payments, perhaps through the use of a discount rate.
- Refunds of overpayments detected through enforcement could be weighted differently from payments of additional assessed tax.
- Private compliance costs could be weighted differently from government costs.

These weights generally cannot be derived theoretically or empirically, and so must be determined by consensus among the ultimate decision-makers, reflecting their best judgments.

Other Constraints: Within a given fiscal year, the IRS operates under a fixed budget and a fixed Internal Revenue Code. However, IRS actions are often constrained by other things, as well, such as the geographic distribution of taxpayers and employees, the job market in each location, employee training needs, and government hiring rules. As long as these constraints can be quantified, they can be accounted for in a model that seeks to maximize net benefits.

Nonrevenue Benefits: Some IRS functions (e.g., most of the Tax-Exempt/Government Entities operating division, much of MITS—Information Technology support—and other support functions) are not directly revenue-related. There are several options for addressing these functions:

- If the benefits of these functions can be expressed in dollars, they can be included among the other benefits and weighted separately.
- If the cost of these functions is mostly overhead to facilitate prefilling, filing, and postfiling operations, then that should be reflected in the cost of those programs.
- Otherwise, resources can be allocated to these functions outside of an optimization framework—much as they are now.

Principle #9. Three forms of fairness are accounted for by the benefits and costs described above:

- Underpayments of tax are prevented or corrected to the maximum extent possible, given the weight chosen for this benefit.
- Overpayments of tax are prevented or corrected to the maximum extent possible, given the weight chosen for this benefit.
- All activities would be funded (and all workload would be selected) on the same objective basis. Taxpayers and taxpayer groups would all be treated similarly because enforcement and service activities would all be guided by this simple objective. This accounts for both direct and indirect effects (not direct alone), and all activities and workload would compete for resources based on their marginal cost-effectiveness in producing net benefits consistent with our ultimate objective.

Principle #10. If we take all of the right benefits and costs into account, then it is not optimal to pursue noncompliance that is not cost-effective at the margin. But is it fair? The answer is both yes and no. It is important to remember that the objective includes as a benefit (a major benefit) the indirect impact of IRS activities on voluntary compliance—both among the taxpayers we contact and those in the general population whom we do not contact. As long as the objective takes into account all of the benefits that taxpayers and the IRS realize from our efforts (as well as all of the costs), there is no remaining fairness issue. For example, pursuing work that is less cost-effective would mean that less of the tax gap would be closed, which would be unfair to compliant taxpayers. Ultimately, the fairness to noncompliant taxpayers (potential targets of enforcement) is not as important as the fairness to compliant taxpayers. However, having said that, we must not ignore pockets of high noncompliance. We should use the tax gap as a guide in allocating our *research* resources, developing better (more cost-effective) ways to improve voluntary compliance and enforcement where the risk is the greatest and the gains are likely to be greatest, as well. (See also Principle #6 above.)

Principle #11. Reduction of taxpayer compliance costs (burden) is not an independent objective. Taxpayer costs are an integral component of the net benefit calculation, but reduction of burden is not a stand-alone objective. In fact, imposing some kinds of burden on taxpayers and third parties is sometimes the best way to maximize net benefits because the alternatives provide far fewer benefits and/or cost much more. Therefore, taxpayer

burden reduction should not be viewed as a goal that is independent of the overall IRS objective. This is easier to understand when we see the distinction between necessary burden and unnecessary burden.

- Necessary burden includes all cost-effective alternatives to more onerous methods (e.g., IRS action alone). That is, private compliance costs are necessary if they produce more net benefit than would be possible without them. A clear example of necessary burden is the cost of existing third-party information reporting, which is far more cost-effective at promoting voluntary compliance than an alternative IRS enforcement program would be in the absence of such reporting. More subtle is the fact that, though IRS forms, instructions, and publications impose some burden on taxpayers, they are far less burdensome than the alternative: an Internal Revenue Code (and regulations) not accompanied at all by forms, instructions, and other information that help taxpayers to understand and comply with the law.
- Unnecessary burden, in contrast, does not increase net benefits cost-effectively. In this case, more cost-effective alternatives should be found. For example, although tax forms and instructions are not inherently unnecessary, any confusion or unnecessary work that they cause is unnecessary burden.

Reducing unnecessary burden and imposing necessary burden are means to the end of increasing net benefits; they are not ends in themselves.

Principle #12. Customer satisfaction and employee satisfaction (as measured by surveys) are also means to achieve our ultimate objective; they are not ends that compete with that objective. We want taxpayers to be satisfied not only because it is right, but because, when they are satisfied, they are more likely to be compliant (increasing the benefits). Likewise, when employees are satisfied, they may work more efficiently, thereby increasing the benefits and/or reducing costs. We would not want to increase customer or employee satisfaction if that meant reducing the net benefit (e.g., by reducing voluntary compliance or by increasing costs). Although tax policy is often formulated to maximize some notion of overall social welfare, once that policy is fixed in the tax laws and regulations, the role of tax administration is to use its resources as cost-effectively as possible to collect as close to the aggregate amount of tax imposed by law as possible. Neither customer nor employee satisfaction is generally a good measure of how well IRS has achieved that objective.

Priorities: How Do We Get From Here to There?

Principle #13. Senior IRS decisionmakers need to develop consensus on the components and form of the ultimate objective (including the weights assigned to the components). These cannot all be derived empirically or theoretically. This consensus could change over time, but it should reflect their best judgment (working in concert with the Department of the Treasury and Congress, as well as with taxpayer and other stakeholder groups) of what the IRS should try to achieve. Furthermore, it will take determined leadership to ensure that all IRS functions cooperate to compile the data necessary to estimate the marginal cost-effectiveness (bang for the buck) of each activity as a function of level of effort, and to update those data annually.

Principle #14. IRS needs to take a long-term view of compiling the right data and developing the estimates necessary to model optimal resource allocation, etc. It may take more than 10 years, but we should be closer to ideal 10 years from now than we are today. It will require a serious cross-functional effort (operations and research organizations working together). (The IRS Enforcement Revenue Information System is an instructive precedent for such an undertaking. It was costly to develop in terms of dollars and time, but it is now relied upon as a crucial IRS database—compiling very useful data on every IRS enforcement case.) It is helpful that the IRS, with the support of Treasury and the Congress, has committed to a long-term National Research Program to collect taxpayer compliance data, which may prove useful in estimating the indirect effects of IRS programs on the voluntary compliance behavior of the general population, but other data will be needed specifically to estimate the marginal cost-effectiveness of those operational programs (see Principle #15).

Principle #15. IRS needs to estimate marginal direct and indirect benefit/cost for each activity (both enforcement and service activities) as a function of resource levels. Short-term approximations can be used and improved over time. Specifically, the IRS research community, working with IRS leadership from all of the divisions and functions, should pursue several critical priorities to move the Service in the right direction:

Short Term (2 Years): Even though good, comprehensive estimates of marginal benefit-to-cost will require much more time to develop, much can be done in the short term. Three initial priorities stand out:

Develop consensus on key components of our ultimate objective:

- What benefits and costs (to IRS and to taxpayers) should be included;

- What relative weights to assign to the various benefits and costs; and
- What discount rate to apply to future amounts.

This consensus is foundational to defining what we should be trying to achieve, and must reflect the judgment and commitments of senior IRS leadership as well as key stakeholders. (See Principle #13.)

Develop preliminary estimates of marginal direct revenue/cost for each program as a function of budget outlay. This could be based on Enforcement Revenue Information System (ERIS) data on total enforcement revenue collected (which would not include all benefits, but would be a step in the right direction) and on budget data on expenditure by program. Such an analysis would not be as detailed as would be possible with data that could be compiled over a longer term, but it would be much better than using average amounts, recommended audit results, and not taking cost-effectiveness into account at all.

Derive assumed marginal indirect revenue/cost curves for each program based on the consensus judgments of senior IRS decisionmakers. This would be an update of a 1998 exercise conducted by PricewaterhouseCoopers in which senior IRS executives developed consensus on the relative magnitudes of the indirect effects of various enforcement and service activities.² These relative magnitudes could then be used to derive presumed indirect revenue/cost curves for each program, based on a curve estimated statistically for Examination.³ The derived curve for each program would be a multiple or fraction of the known curve, where the multiples and fractions are based on the consensus relative magnitudes.

Although this approach is far from ideal, it does have some advantages. It can move us in the right direction in a relatively short time. It captures the best judgment of key decisionmakers outside of budget pressures, replacing much of the subjective and inconsistent reasoning that is often applied during the budget process. It also allows resources to be allocated according to consensus rules based on the right principles, and can be used in conjunc-

² See Plumley (2002).

³ See, for example, Plumley (1996).

tion with estimates of the marginal direct effect, which would be estimated separately. Fundamentally, this approach is better than current attempts to foster voluntary compliance by maintaining arbitrary minimum coverage constraints, and it is certainly better than doing nothing at all to account for indirect effects.

Long Term (10 Years): The best way to develop robust and comprehensive estimates of marginal benefit-to-cost for each IRS activity is to compile the right kind of data and to apply appropriate research methods. Both of these will require a long-term effort and will need to be updated regularly. (See Principle #14.) This research should include several key components:

Develop consensus among IRS researchers and a panel of academic researchers as to the types of methodologies likely to be able to produce good estimated marginal benefit/cost curves as a function of level of effort (or cost) for each IRS activity (both enforcement and nonenforcement activities). The long-term data needs should flow from the selected methodologies; we should not expect to be able to estimate marginal benefit/cost functions solely from data systems designed to manage operational programs.

Compile detailed cost data for each activity. This will be needed for any estimation method but should be compiled in a way that will facilitate the selected approach. For example, it will have to be decided what variables the costs should be associated with for the eventual analysis (e.g., type of activity, type of taxpayer, geographic area, etc.). The cost data should be comprehensive (capturing both hours and dollars), including direct time applied, indirect time (training, leave, recordkeeping, etc.), support costs (human resources, management, secretarial, etc.), capital and overhead costs (equipment, supplies, facilities, etc.), and contract costs (for data, services, etc.). Compiling cost data of this nature in a consistent way across all IRS activities will be a major cross-functional undertaking—similar, perhaps, to what was needed to create the Enforcement Revenue Information System (ERIS). It likely would not impose more administrative requirements on front-line employees, but instead would capture existing information in a common format. Like ERIS, it would be hard to create but would become an invaluable source of information once it is fully functional.

Compile detailed and standardized output data for each activity. Much of this is probably already compiled, but not in a standardized way across all activities (e.g., by type of taxpayer, geographic area, etc.). This would include direct enforcement revenue (for estimating marginal direct effects) but also outputs from service activities.

Compile detailed data on taxpayer compliance behavior, such as returns filed (and when) and amounts of tax paid voluntarily and on time (including the method and timing of payment), as well as late payments (particularly nonenforced late payments)—all by type of tax and the same geographic and type of taxpayer indicators defined for the cost and output data. These data would be used to develop estimates of the indirect effects of both enforcement and service activities—estimates that will likely require on the order of 10 years of data.

Compile data on non-IRS factors likely to influence compliance behavior, such as economic, demographic, and attitudinal variables, as well as tax law changes. These would need to be compiled in the same geographic and time dimensions used for compiling all the data described above.

Introduce more variation (geographically and over time) in spending on specific IRS enforcement and nonenforcement activities than would normally be the case. In essence, these variations would be controlled field experiments designed to estimate the marginal impact of these activities on taxpayer compliance behavior, where all other relevant factors (both IRS and non-IRS factors) are controlled for statistically (rather than with standard control groups).⁴

⁴ Multiple factors undoubtedly shape voluntary compliance behavior: presumably most IRS activities plus other factors outside of IRS control. A change in taxpayer behavior that appears to be associated with a change in IRS actions could actually be caused by one or more other factors, so that we must control for all relevant factors simultaneously. A simple way to do this is to identify a control group that is identical to the treatment group in all respects other than the treatment. However, given the large number and wide variety of potential determinants of compliance behavior (and the fact that we do not necessarily know what they all are), a control group in this context would probably have to be quite large and dispersed across the country, raising the likelihood that it would be influenced by the treatment. Furthermore, if a significant determinant of compliance behavior (e.g., religion) is not controlled for, its impact would be attributed incorrectly to the IRS intervention. That would not be a problem when controlling for compliance determinants statistically—as long as the excluded variable was not correlated with the IRS variables of interest.

Conclusion

IRS can undoubtedly make better use of existing and new resources to achieve greater benefit and impose less cost on taxpayers. To do that, we need to be explicit about our ultimate objective and allocate resources at the margin accordingly. Developing that capability will likely take a concerted, long-term (10-year), cross-functional effort to gather the right data every year and to introduce extra variation in our activities, making it easier to estimate their direct and indirect impacts at the margin. We can take other steps in the short term to make some improvements. However, allocating resources according to the distribution of the tax gap is not the right approach.

Acknowledgments

I wish to thank Joel Slemrod for comments and suggestions on earlier drafts of this paper, and for bearing with me on points of (what I believe to be minor) disagreement. Janet McCubbin and Mark Mazur also provided valuable suggestions.

References

- Plumley, Alan H. (1996), *The Determinants of Individual Income Tax Compliance: Estimating the Impacts of Tax Policy, Enforcement, and IRS Responsiveness*, IRS Publication 1916 (Revised 11-96).
- Plumley, Alan H. (2002), "The Impact of the IRS on Voluntary Tax Compliance: Preliminary Empirical Results," National Tax Association, *Proceedings of the Ninety-Fifth Annual Conference*, pp. 355-363.
- Plumley, Alan H. and C. Eugene Steuerle (2004), "Ultimate Objectives for the IRS: Balancing Revenue and Service," in *The Crisis in Tax Administration*, edited by Henry J. Aaron and Joel Slemrod, Brookings Institution Press, Washington, D.C., pp. 311-346.